Admissibility Closure

Deriving Quantum Structure from Finite Distinguishability and Irreversible Commitment

Keith Taylor *VERSF Theoretical Physics Programme*

Plain-language summary

Quantum mechanics is usually taught as a list of rules you have to accept: there is a thing called a wavefunction, it lives in a complex vector space, probabilities are obtained by squaring its size, and measurement collapses it. The rules work spectacularly well, but they are presented without explanation. *Why* a complex space rather than a real one? *Why* squaring rather than some other operation? *Why* a special "measurement" rule that doesn't apply to anything else?

This paper argues that those rules are not independent assumptions. They are the unavoidable consequences of two simple physical facts plus two consistency requirements that any universe capable of containing stable, irreversible records — facts that stay fixed once they happen — would have to satisfy.

The two physical facts are these. First, **finite distinguishability**: with finite physical resources you can only tell finitely many things apart at once. Second, **irreversible commitment**: some processes produce records that cannot be undone — this is what we mean by a measurement actually happening, and it is what gives time its direction. The two consistency requirements are that **independent systems combine in the natural way** (the joint description of two unrelated things is built from their separate descriptions) and that **physical structure does not depend on the labels we use** to describe it.

From these four ingredients alone, the paper shows, the entire kinematic skeleton of quantum mechanics follows. There must be a finite set of alternatives prior to any measurement. Those alternatives must be represented by complex numbers — not real numbers, not quaternions, the choice is forced by the consistency requirements. The probability of any outcome must be the squared size of the corresponding amplitude — the Born rule — because no other rule is compatible with the four ingredients. And measurement itself is no longer mysterious: it is a competitive race between different outcomes to produce the first irreversible record, and the rule for which one wins is fixed by the kinematic skeleton, not added on top.

The headline claim is that quantum mechanics is not an arbitrary set of postulates. It is the only kinematic architecture compatible with the existence of stable facts. Any universe in which records can be made and remembered must be quantum-mechanical at this level of description.

Abstract

We present a unifying synthesis of the VERSF programme by showing that the principal structures of quantum mechanics — Hilbert-space representation, reversible pre-measurement evolution, irreversible measurement, and the Born rule — arise as **forced** consequences of two primitive physical constraints (finite distinguishability and irreversible commitment) together with two consistency requirements (compositional consistency and observer invariance), the latter unpacked into its topological and one-parameter components.

From these we derive in sequence: (i) a finite, structured pre-commitment domain; (ii) a reversible group action on the *set* of pre-commitment alternatives, separated explicitly from the reversible group action on their *representation*, with the additional premise (P) — that distinguishability change occurs only through the (A2) commitment channel — made fully explicit; (iii) an amplitude representation over a continuous division algebra carrying a U(1) phase, with each step (linearity, finite-dimensionality, associativity, division-algebra closure) named and cited; (iv) a bilinear, gauge-invariant probability functional; (v) a rank-one reduction forcing $p = |\psi|^2$ jointly with bilinearity, derived from compositional consistency in a single argument whose explicit calculation is carried out in Appendix A; (vi) measurement as a first-passage commitment process on the underlying substrate, with an explicit characterisation of which substrate rate laws are admissible.

We position the result against existing axiomatisations (Hardy 2001, Masanes–Müller 2011, Chiribella–D'Ariano–Perinotti 2011, Gleason 1957, Busch 2003, Zurek 2005) and identify what is distinctive about the VERSF route: not the existence of a derivation, but the substrate-grounded measurement account that makes the Born rule a *consequence* of admissible dynamics rather than a probabilistic postulate. The closure theorem of §9 is calibrated accordingly and split into a kinematic part (forced by the primitives) and a substrate-dependent part (forced by the primitives plus the Tick–Bit dynamics, modulo a stated consistency requirement).

Table of Contents

---

1. Introduction

Quantum mechanics is conventionally introduced through a list of postulates — a complex Hilbert space, unitary evolution, self-adjoint observables, the Born rule, and a projection postulate for measurement. The structure works, but it is *posited*: the postulates do not explain why these particular structures, rather than some internally consistent alternative, describe physical systems.

The VERSF programme has produced several apparently independent recoveries of central quantum structures. The Born rule alone has been obtained through:

- distinguishability geometry on the pre-commitment domain (Fubini–Study route);
- conservation of an information-flux current (Noether route);
- a first-passage analysis on the Tick–Bit substrate.

That three routes converge on the same law is suggestive. The natural question is whether they are independent results or projections of a single constraint.

This paper argues for the second reading. We identify the constraint — **admissibility** under finite distinguishability and irreversible commitment — and show that each route is a face of it. The Born rule is the unique probability assignment compatible with admissibility itself, and quantum kinematics is the unique architecture compatible with admissibility modulo a single, named field-selection step.

1.1 Relation to existing axiomatic programmes

The conclusion that quantum mechanics is the unique theory satisfying a small set of operational or informational axioms is not new. Hardy (2001) derives the quantum formalism from five "reasonable" axioms; Masanes and Müller (2011) sharpen this to four; the Chiribella–D'Ariano–Perinotti (2011) programme derives QM from informational principles including purification; Gleason (1957) and Busch (2003) establish the Born rule from non-contextuality of probability measures on lattices/effects; Zurek (2005) derives the Born rule from environment-assisted invariance under unitary symmetry.

The present paper is *not* in competition with these. The (A1)–(A4) primitives we use overlap meaningfully with theirs: compositional consistency is close to tomographic locality, observer invariance is close to continuous reversibility, finite distinguishability has analogues in the finite-information assumptions of Hardy and CD'AP. What is distinctive in the VERSF route is twofold:

1. The primitives are framed in terms of *records* and *commitment* rather than operational probabilities, so the derivation is naturally embedded in a substrate-level account of what measurement physically *is*.
2. The first-passage account of measurement (§7) supplies a dynamical mechanism that the operational programmes leave open — they derive the Born rule but not its physical realisation.

The convergence of multiple distinct axiom sets on the same structure is itself a finding worth marking: it suggests that the destination is robust under reformulation, and that any one set of primitives — including ours — is one route among several to a structurally inevitable conclusion.

1.2 A note on the term "admissible"

We use "admissible" throughout in one specific sense: **compatible with primitives (A1)–(A4) and the bridging premise (P) of §2**. A representation, a transformation, a probability assignment, or a dynamics is admissible if it is consistent with finite distinguishability, irreversible commitment, the (P) channel premise, compositional consistency, and observer invariance. This is the only sense in play; the word is not used to mean "physically reasonable" or "well-defined" in any broader register.

The argument is presented in seven stages: (§2) the primitive admissibility constraints; (§3) the necessity of a reversible pre-commitment structure, with representation and ontic levels separated; (§4) the emergence of amplitude geometry, with each step in the field-selection chain made explicit; (§5) the unification of the Born-rule routes, with their relative logical status clarified; (§6) a uniqueness theorem; (§7) measurement as physical commitment, with explicit substrate-consistency conditions; (§8) time and irreversibility (forward pointer to companion papers); (§9) the closure theorem, calibrated to what the body of the paper establishes.

---

2. Minimal Admissibility Framework

We adopt two primitive physical constraints, two consistency requirements, and one bridging premise.

(A1) Finite Distinguishability

Within any bounded resource budget, only finitely many distinctions can be physically resolved.

Equivalently: every physical observation produces a record drawn from a *finite* alphabet, and no physical procedure resolves an unbounded number of distinctions in finite resources. Earlier VERSF work establishes that infinite distinguishability is incompatible with the existence of stable irreversible records.

(A2) Irreversible Commitment

There exist physical processes that map several admissible prior states onto a single posterior state with no admissible inverse.

Such processes constitute *fact production*. They are the operational content of the second law: they ground the entropy gradient and the temporal ordering of records.

(P) Channel Uniqueness (bridging premise)

Distinguishability change in admissible dynamics occurs *only* through the (A2) commitment channel. No admissible transformation merges or creates distinctions in $\mathcal{A}$ except as an instance of (A2).

(P) is the premise that converts (A2) from "commitments have this form" into "transformations of this form are commitments." In earlier presentations this step was implicit, leaving Proposition 3.2 ambiguously placed between a derivation and a definitional gloss. We promote (P) to a named premise, so that §3.2 is genuinely a non-trivial proposition. (P) is defensible on substrate-level grounds — a commitment-free interval has no admissible mechanism for merging or creating distinctions, since records are produced only by commitments — but we treat it here as a stated premise rather than a derived result, and flag it as such.

(A3) Compositional Consistency

When two systems are treated independently, their joint distinguishability structure is the product of their separate structures. No distinctions are created or destroyed by the act of composing, and the joint-system probability assignment is determined by the marginal assignments together with their correlations.

This is close to what the operational literature calls *tomographic locality*: the state of a composite system is determined by the joint statistics of local measurements. The strength of (A3) as stated here — full tomographic locality, not a weaker variant — matches that used in operational reconstructions of quantum theory (Hardy 2001, Masanes–Müller 2011, Chiribella–D'Ariano–Perinotti 2011). Weakening (A3) — for instance, admitting hyper-decoherent theories, non-locally-tomographic theories, or generalised probabilistic theories without strong product structure — admits broader classes of admissible architectures that the present argument does not exclude. We adopt the strong form because it is what the argument requires; the dependency is acknowledged here at the primitive level so that the conditional status of the closure claim is visible from the outset rather than only at the points (Theorems 4.1, 4.2, 4.3) where (A3) does its work.

(A4) Observer Invariance

Physical distinctions are independent of the labels or representational basis used to describe them. Any admissible structure must be invariant under continuous re-parametrisation of the descriptive frame.

The "continuous" qualifier in (A4) is doing substantial work later in the paper, and it is really two distinct claims bundled together. We unpack them:

- **(A4a) Topological reversibility.** Admissible representational transformations form a topological group rather than a discrete one.
- **(A4b) One-parameter coverage.** That topological group admits continuous one-parameter subgroups acting non-trivially on each component of a representation.

(A4a) enters Proposition 3.3 and Theorem 4.1 Step 1; (A4b) enters Theorem 4.1 Step 2 and the field-selection argument of Theorem 4.2. A reader sympathetic to the framework but skeptical of (A4b) — for instance, one interested in fundamentally discrete substrates — should note that the arguments dependent on (A4b) can fail without (A4a) failing. We mark this throughout.

These primitives plus (P) define **physical admissibility**: the minimal conditions a substrate must satisfy to support a universe in which stable, irreversible records exist.

We make no further postulate. We do not assume linearity, Hilbert space, complex numbers, unitarity, or the Born rule. Each is derived (or, in the case of the field, derived modulo the field-selection step we make explicit).

---

3. The Pre-Commitment Domain: Two Levels

Section 3 has been the site of a conflation in earlier presentations between two distinct things: transformations of the *set* of alternatives (ontic level) and transformations of the *representation* of those alternatives (representational level). Both are reversible in admissible dynamics, but for different reasons. We separate them.

Proposition 3.1 — Pre-commitment alternatives must exist

A commitment event resolves alternatives. If no distinguishable alternatives existed prior to the event, "commitment" would have no operational content.

Proof. By (A2), commitment maps multiple priors to one posterior. The cardinality of the prior set must be ≥ 2; otherwise the map is the identity and no fact is produced. By (A1), this prior set is finite. Hence the substrate must support, prior to any commitment, a finite collection $\mathcal{A} = \{a_1, \ldots, a\_N\}$ of distinguishable alternatives. ■

Proposition 3.2 — Pre-commitment dynamics preserve the alternative set

Let T be a dynamical transformation acting on the substrate during a pre-commitment interval — that is, an interval during which (A2) does not fire. Then under (P), T preserves the set $\mathcal{A}$ of admissible alternatives: it neither merges nor creates distinctions.

Proof. Suppose T merges two distinct alternatives a_i, $a_j \in \mathcal{A}$ into one. By (P), distinguishability merging in admissible dynamics occurs only through the (A2) commitment channel. But T was assumed to act in a pre-commitment interval, in which no commitment event fires. Contradiction.

Suppose T creates a distinction not previously in $\mathcal{A}$. By (P), distinguishability creation in admissible dynamics occurs only through (A2). Again contradiction. ■

The proposition is now a genuine consequence: it requires both the assumption of a pre-commitment interval *and* the channel-uniqueness premise (P). Stripping (P) leaves Proposition 3.2 as a *definition* of what counts as a pre-commitment interval, which is also coherent but does different work in the argument. We adopt the (P)-premised reading throughout because it permits §4 to treat pre-commitment intervals as a priori populated rather than vacuously definable.

This establishes set-level reversibility: pre-commitment dynamics act as bijections on $\mathcal{A}$.

Proposition 3.3 — Representation-level transformations form a continuous group

The transformations of the representation of $\mathcal{A}$ — i.e. changes of basis, relabelling, and continuous reparametrisations of how each alternative is described — form a continuous group of bijections that acts trivially on the set $\mathcal{A}$ itself.

Proof. By (A4), no descriptive label is privileged. The full set of admissible representations is therefore an orbit under a relabelling group G. Two admissible representations connected by G describe the *same* underlying alternative set $\mathcal{A}$ (no distinctions are created or merged by relabelling), so G acts trivially on $\mathcal{A}$. By (A4a), G is a topological group. ■

Why the separation matters

Proposition 3.2 establishes that pre-commitment dynamics are reversible at the level of the set of alternatives. Proposition 3.3 establishes that there is a continuous symmetry group acting at the level of how those alternatives are represented. These are different statements with different premises. Earlier formulations collapsed them, which left the impression that *any* non-bijective transformation was automatically a commitment — which would rule out, for example, a perfectly admissible change of basis that maps one orthonormal frame on $\mathcal{A}$ to another. The corrected version permits exactly the structure we need in §4: a fixed alternative set carrying a continuous group of representational transformations.

Interpretation

The composite of Propositions 3.1–3.3 is what is conventionally called *quantum superposition*: a finite alternative set together with a continuous group of representation-level transformations that preserve the set. In the standard formulation this is postulated; here it is forced.

---

4. Emergence of Amplitude Structure

We now ask what mathematical structure is required to *represent* the pre-commitment domain established in §3. The argument proceeds in three theorems, with each step in the chain (linearity, finite-dimensionality, associativity, division-algebra closure, field selection) named.

Theorem 4.1 — Amplitude representation over a continuous division algebra

The reversible representation-level transformations on $\mathcal{A}$ require a representation of each alternative as a vector component over a finite-dimensional, continuous, associative division algebra $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}, \mathbb{H}\}$, with a continuous compact group acting on each component.

Proof.

Step 0 — Linearity. Before assigning components, we must show that the representation is *vectorial* rather than the more general convex state space of the GPT (generalised probabilistic theory) literature. Compositional consistency (A3) plus continuity (A4a) constrains the joint state space of independent subsystems to be determined by marginal statistics and continuously parametrised correlations. Hardy (2001 §6) and Masanes–Müller (2011 §III) show that under finite distinguishability plus continuous reversibility plus tomographic locality, the state space is forced to be a linear cone closed under tensor product; this is the content of the "simplicity," "subspace," and "continuity" axioms in those works. We adopt this as established in the operational literature and import it: the admissible representation is linear.

It is worth being explicit that this *does not* dismiss GPTs en bloc. Many GPTs (Barrett's no-signalling polytope, classical theory, real quantum theory, etc.) admit well-defined tensor products and satisfy weak forms of tomographic locality. What the cited results show is that the *joint* satisfaction of (A3) in the strong form, (A4a), and finite distinguishability picks out *linear* state spaces from this broader class. The operational details of this filtration are in the cited papers; we take them as input. **The conclusion is therefore conditional on the strong form of compositional consistency (A3) together with continuous reversibility (A4a); weaker forms — for instance, admitting non-locally-tomographic theories or relaxing continuity — admit broader GPT classes that the present argument does not exclude.** We mark this dependency rather than gloss it: the linearity step inherits whatever conditional status its source axioms carry.

Step 1 — Carrier is a finite-dimensional associative division algebra over $\mathbb{R}$. By Proposition 3.1, $\mathcal{A}$ has finite cardinality N, so an admissible linear representation assigns to each alternative a component value drawn from some carrier algebra $\mathbb{F}$ over $\mathbb{R}$. Four properties of $\mathbb{F}$ must be argued:

(a) **Finite dimensionality over $\mathbb{R}$.** N is finite, the representation is N-dimensional over $\mathbb{F}$, and physical admissibility on a finite alternative set forbids unbounded scalar resources per component (this is a finite-distinguishability consequence: an infinite-dimensional carrier would resolve infinitely many distinctions per component, violating (A1) on the joint representation).

(b) **Associativity.** By Proposition 3.3, G is a group, and the action of G on linear representations composes associatively; for the representation of composite systems via tensor product (A3), the underlying scalar algebra must support associative multiplication so that $(g_1g_2)\cdot(g_3g_4) = g_1\cdot(g_2g_3)\cdot g_4$ is well-defined on tensor products. Non-associative algebras (notably the octonions $\mathbb{O}$) cannot support tensor-product representations of associative groups in the required way; this is a standard result in representation theory (e.g. Baez 2002 §3.4). Octonions are accordingly excluded.

(c) **Continuity.** By (A4a), G is a topological group acting continuously on representations, so the algebra operations on $\mathbb{F}$ must be continuous in the standard topology on $\mathbb{R}^n$.

(d) **Division-algebra closure.** This is the step that earlier presentations did not argue explicitly. We restrict, by construction, to *one-dimensional* $\mathbb{F}$-modules per alternative — each $a_i \in \mathcal{A}$ carries a single component $c_i \in \mathbb{F}$ — so that the representation-group action on a component reduces to scalar multiplication on a one-dimensional $\mathbb{F}$-module. The restriction is forced by (A1): a higher-dimensional internal structure on each alternative would resolve sub-distinctions within a_i, splitting it into multiple admissible alternatives, which contradicts the construction of $\mathcal{A}$ as the *finest* finite alternative set produced by Proposition 3.1. With this restriction in place, the argument runs as follows. By Proposition 3.2, pre-commitment dynamics are reversible bijections at the level of $\mathcal{A}$. By Proposition 3.3, the representation group G acts reversibly on the representation; so for every $g \in G$ and every component $c \in \mathbb{F}$, the action $g \cdot c$ must have an inverse $g^{-1} \cdot (g \cdot c) = c$. For the action to be linear and reversible at the component level — i.e. for the component-level action of any g to be an invertible element of $\mathbb{F}$ acting by multiplication on a one-dimensional $\mathbb{F}$-module — every non-zero element of $\mathbb{F}$ must possess a multiplicative inverse. (Equivalently: a continuous, faithful, linear action of a topological group G on a one-dimensional 𝔽-module is invertible only when 𝔽 is a division algebra.) Hence 𝔽 has no zero divisors and every non-zero element is invertible: 𝔽 is a division algebra.

Combining (a)–(d), 𝔽 is a finite-dimensional, associative, continuous division algebra over ℝ. By the Frobenius theorem (Frobenius 1878), the only such algebras are ℝ, ℂ, and ℍ.

Step 2 — Per-alternative continuous symmetry. It remains to show that G acts non-trivially on each individual component, not merely as a global relabelling of the index set. This is the step where (A4b) enters and where superposition (§3) does work. Coherent combinations of alternatives — admissible by Proposition 3.2 because the set $\mathcal{A}$ is preserved — must themselves be admissibly representable. That is, for any admissible coefficients $(c_1, \ldots, c_N)$, the combination $\Sigma_i c_i a_i$ must transform admissibly under G. If G acted only as a permutation of the index i, it would not act on the coefficients c_i at all, and continuous reparametrisations of those coefficients (admitted by (A4b)) would be unrepresentable. Hence G must include a continuous compact action on each component c_i.

The "minimal" such action — by which we mean the smallest closed connected subgroup acting faithfully and continuously on a single component — is determined by the algebra:

- **𝔽 = ℝ:** the smallest closed connected continuous subgroup of GL(1, ℝ) is trivial. The largest discrete faithful subgroup is $\mathbb{Z}_2$ (sign flip), but it is not continuous.
- **𝔽 = ℂ:** the smallest closed connected continuous subgroup acting faithfully on a single component is $U(1) \cong S^1$.
- **𝔽 = ℍ:** the smallest closed connected continuous subgroup acting faithfully on a single quaternionic component is $SU(2) \cong S^3$ (one-sided action; the two-sided action gives Spin(4)).

The ℝ entry is highlighted: real representations admit *no* continuous one-parameter group acting on a single component, because GL(1, ℝ) is itself discrete in its connected component (the positive reals act by scaling, but rescaling changes $|c|$ and so changes distinguishability — which violates (A4) since the underlying alternative set is unchanged). This is the precise sense in which (A4b) discriminates ℝ from ℂ and ℍ; we return to it in Theorem 4.2. ■

Theorem 4.2 — Selection of ℂ over ℝ and ℍ

Among the three division algebras admitted by Theorem 4.1, only ℂ supports a representation that simultaneously satisfies (A3) tomographic locality and (A4b) per-component continuous symmetry for all $N \geq 2$.

Proof.

ℝ *fails on per-component continuous symmetry.* This is the precise version of the argument hinted at in earlier drafts. Real quantum theory (in the Stueckelberg formulation) has a perfectly continuous orthogonal symmetry O(N) acting on the *global* representation, so it would be wrong to say "real quantum theory has no continuous symmetries." The relevant claim is sharper: real quantum theory admits no continuous one-parameter group of *per-component phase* transformations, because GL(1, $\mathbb{R}$) has no non-trivial connected subgroup that preserves the squared modulus $|c|^2$. Since the bilinear/rank-one functional of Theorem 4.3 requires a per-component U(1) phase symmetry to suppress higher-order terms in the (n,n) expansion, and real representations cannot supply this, the §4.3 derivation collapses over $\mathbb{R}$. Equivalently: the rank-one structure of probability functionals over $\mathbb{R}$ is not forced — additional families of admissible quadratic functionals exist that are not of the form $|\langle\varphi|\psi\rangle|^2$, and (A3) factorisation alone is not strong enough to exclude them. The continuous per-component phase symmetry is what closes the rank-one argument, and $\mathbb{R}$ does not provide it.

$\mathbb{H}$ fails on tomographic locality. Over $\mathbb{H}$, the joint state space of two systems is *not* determined by the marginal state spaces and their correlations: quaternionic tensor products carry additional non-commutative phase information that has no marginal correlate. This is the well-known failure of tomographic locality for quaternionic quantum mechanics (Araki 1980; Wootters 1990; Baez 2012). It directly violates (A3).

$\mathbb{C}$ satisfies both. Complex quantum theory has continuous U(1) per-component symmetry (satisfying (A4b)) and is tomographically local (satisfying (A3)). It is the unique division algebra in the Frobenius list passing both filters.

Hence

$\psi \in \mathbb{C}^N$, $\psi = (\psi_1, \ldots, \psi_N)$, with global U(1) gauge $\psi \mapsto e^{i\theta} \psi$. ■

The field-selection step is now fully explicit, with the $\mathbb{R}$-exclusion argument given in the form that survives the obvious counterexample (real QT has continuous O(N) symmetry).

Scope of the field-selection step

A note on what this argument does and does not establish. The selection of $\mathbb{C}$ is contingent on (i) adopting tomographic locality in its strong form, as encoded in (A3), and (ii) requiring per-component continuous phase symmetry, as encoded in (A4b). Weakening (i) — admitting hyper-decoherent or non-locally-tomographic theories — would re-open the door to quaternionic or hybrid alternatives. Weakening (ii) — admitting fundamentally discrete substrates — would re-open the door to real quantum theory and to broader operational frameworks. The closure claim of this paper is therefore conditional on the strong reading of (A3) and on (A4b); we mark this contingency rather than gloss it. The field-selection step is the only place in the kinematic derivation where genuinely substantive operational axioms do the discriminating work, and the argument's force depends on accepting them.

Theorem 4.3 — Bilinearity and rank-one factorisation, derived jointly

Any admissible probability functional $P(a \mid \psi)$ is bilinear in (ψ, ψ^*) with rank-one Hermitian kernel:

$P(a \mid \psi) = |\langle\varphi(a) \mid \psi\rangle|^2$.

Proof sketch. We establish the bilinear form and the rank-one structure in a single argument; the explicit calculation is given in Appendix A.

Setup. P is a real-valued continuous function of $\psi \in \mathbb{C}^N$. By (A4) phase invariance, $P(e^{\wedge}(i\theta)\ \psi) = P(\psi)$ for all θ. By the existence of a real-valued probability, P is invariant under joint complex conjugation of its arguments. We seek the most general such functional consistent with (A3).

Compositional constraint. Consider two non-interacting subsystems with joint amplitude $\psi = \psi^{(1)} \otimes \psi^{(2)}$. By (A3), the joint probability of independent outcomes (a, b) must factorise:

$P(a, b \mid \psi^{(1)} \otimes \psi^{(2)}) = P(a \mid \psi^{(1)}) \cdot P(b \mid \psi^{(2)})$.

This must hold **for all admissible product states** $\psi^{(1)} \otimes \psi^{(2)}$, with arbitrary relative norm and arbitrary component structure on each factor. The universality of this requirement — not just its holding in special cases — is what forces the global suppression of higher-order terms.

Joint derivation. Expand P as a power series in (ψ, ψ^*) compatible with (A4) phase invariance — only joint terms $\psi_i\ \psi_j^*$ survive each level of the expansion. Phase-invariant terms occur at orders (n, n) with $n \geq 1$. Apply the factorisation requirement order by order. At order (1,1) the requirement is that the bilinear form $K^{(12)}$ on the joint space be a tensor product of marginal bilinear forms — a standard tensor-factorisation statement. At order (2,2) and higher, the universal factorisation requirement forces the corresponding kernels to vanish identically: a non-zero (n,n) kernel with $n \geq 2$ contributes terms of total degree 2n in the marginal amplitudes, and matching these against products of marginal expansions across *all* admissible product norms reduces to a Cauchy-style multiplicativity constraint that has only the power-law solution, with the power fixed to 1 by normalisation (Appendix A).

The conclusion is that P is purely first-order:

$P(a \mid \psi) = \Sigma_{ij}\ \psi_i\ K_{ij}(a)\ \psi_j^*$,

with K(a) Hermitian by reality. Re-applying the factorisation requirement at this surviving order forces the joint kernel to be a tensor product of marginal kernels:

$K(a, b) = K^{(1)}(a) \otimes K^{(2)}(b)$.

The Hermitian operators closed under tensor factorisation in this way — i.e. those for which every product-state expectation factorises — are precisely the rank-one positive operators (Appendix A). Hence

$K(a) = |\varphi(a)\rangle\langle\varphi(a)|$,

and substituting,

$P(a \mid \psi) = |\langle\varphi(a) \mid \psi\rangle|^2$. ■

This is the Born rule, derived via a single application of (A3) rather than two overlapping applications. The bilinearity-vanishing and rank-one steps are both consequences of universality over admissible product states, not of independent operational axioms.

---

5. The Born Rule as an Admissibility Fixed Point

We can now exhibit the apparently independent VERSF derivations of the Born rule as projections of the §4 result, and we are explicit about which routes are *internal cross-sections* of §4 and which require additional structure.

5.1 Geometric route (internal cross-section)

Distinguishability symmetry plus phase covariance — both established in §4 — pick out the Fubini–Study metric on $\mathbb{CP}^{N-1}$. The unique invariant probability measure on rays is $p \propto |\psi|^2$. This route is a geometric repackaging of Theorem 4.3; its premises are not independent of §4.

5.2 Informational route (internal cross-section)

The conserved current associated by Noether's theorem to the U(1) symmetry of Theorem 4.1 is

$$J^u = (i/2)(\psi^* \partial^u \psi - \psi \partial^u \psi^*),$$

with density $J^0 = |\psi|^2$. This is the unique current bilinear in ψ that is conserved under admissible dynamics. The premises — U(1) symmetry, bilinearity — are again drawn from §4.

5.3 Envariance route (independent route, with a noted dependency)

This route, due in the operational literature to Zurek (2005), can be reconstructed inside the admissibility framework as follows. Consider a system S entangled with an environment E in the Schmidt form

$$|\psi\_SE\rangle = \Sigma_i c_i |s_i\rangle \otimes |\varepsilon_i\rangle.$$

A swap on S that exchanges $|s_j\rangle \leftrightarrow |s\_k\rangle$ with phases must be undoable by a counter-swap on E if (A2) is to be respected — the joint state must be returned to the same admissible class. This *envariance* of the joint state under correlated swaps forces equal probabilities for outcomes with equal-modulus amplitudes, and a uniform-prior argument over re-encodings then forces $p\_a \propto |c\_a|^2$ for arbitrary amplitudes.

Caveat. Envariance arguments in this form do not derive the entire kinematic structure; they assume the tensor-product Hilbert space, the Schmidt decomposition, and the existence of a unitary group acting on each factor. As Schlosshauer and Fine (2005) and Mohrhoff (2004) note, this presupposes much of the structure that operational programmes (Hardy 2001, MM 2011) work to derive. So the envariance route is independent of *§4 of this paper* in its premises, but not independent of the broader Hilbert-space scaffolding. Its convergence on the same numerical answer as Theorem 4.3 is the relevant point: two derivations starting from different premises agree on the probability law.

5.4 First-passage (Tick–Bit) route (substrate-dependent)

In the Tick–Bit substrate, branches compete to produce the first irreversible record. The first-passage rate to outcome a is $\lambda_a \propto |\psi_a|^2$, so

$P(a) = \lambda_a / \Sigma_b \lambda_b = |\psi_a|^2 / \Sigma_b |\psi_b|^2$.

This route is *not* an internal cross-section of §4. It imports a dynamical posit — the rate law on the Tick–Bit substrate — that is not derivable from (A1)–(A4) alone. Its agreement with §4.3 is therefore a non-trivial *consistency* between the kinematic admissibility structure and the substrate dynamics. That consistency is what licenses §7.

5.5 What the unification claim actually says

Summary. All routes agree on the same kinematic structure; only the substrate route addresses its physical realisation.

The honest reading is the following:

- §5.1 and §5.2 are different geometric and dynamical projections of the §4 result; they do not constitute additional evidence so much as additional *forms* of the same claim.
- §5.3 (envariance) is independent in its premises of §4, modulo the Hilbert-space scaffolding noted above; it converges on the same answer.
- §5.4 (first-passage) is independent in a stronger sense — it sits at the substrate level — and its agreement with §4.3 is a *constraint that admissible substrates must satisfy*.

So the unification is real but more nuanced than "four routes converge." There is one kinematic destination (§4.3) with two natural geometric/dynamical reformulations (§5.1, §5.2), one independent operational route (§5.3) with its own load-bearing assumptions, and one substrate-level realisation (§5.4) that admissible substrates are required to match.

---

6. Uniqueness Theorem

We now state the result that organises §3–4.

Theorem 6.1 — Admissible Probability Uniqueness

Let a physical theory satisfy (A1) finite distinguishability, (A2) irreversible commitment, (P) channel uniqueness, (A3) compositional consistency, and (A4) observer invariance (with topological component (A4a) and one-parameter component (A4b)). Then the only admissible probability assignment over distinguishable outcomes is

$$\mathbf{p_i = |\psi_i|^2 / \Sigma_j |\psi_j|^2},$$

where $\psi \in \mathbb{C}^N$ *is the amplitude representation of the pre-commitment domain.*

Proof structure.

Step	From	Yields	Status
1	(A1)	finite alternative set $\mathcal{A}$, $|\mathcal{A}| = N < \infty$	Prop 3.1
2	(A2) + (P)	set-level reversibility on $\mathcal{A}$	Prop 3.2
3	(A4a)	continuous representation group G	Prop 3.3
4	(A3) + (A4a) + finite distinguishability	linear representation (Hardy 2001, MM 2011)	Thm 4.1 Step 0
5	finite-dim, associative, continuous, division-closed	$\mathbb{F} \in \{\mathbb{R}, \mathbb{C}, \mathbb{H}\}$ (Frobenius 1878)	Thm 4.1 Step 1
6	(A4b) + Prop 3.2	per-component continuous symmetry on $\mathbb{F}$	Thm 4.1 Step 2
7	(A3) tomographic locality + (A4b)	$\mathbb{F} = \mathbb{C}$, U(1) per component	Thm 4.2
8	(A3) factorisation universality + (A4) phase invariance	bilinear rank-one $P = |\langle\varphi|\psi\rangle|^2$	Thm 4.3 + Appendix A

Each row cites the result that establishes it. The chain of inferences (A1) → (A2,P) → (A4a) → (A3+A4a+finite) → Frobenius → (A4b+Prop 3.2) → (A3+A4b) → (A3 univ + A4) is now traceable end-to-end without sketched steps. ■

7. Measurement as Physical Commitment

The closure result of §3–6 is *kinematic*: it establishes the structure of the pre-commitment domain and the form of the probability functional. To turn this into an account of measurement, we need a *dynamics* that produces irreversible records. That dynamics is supplied by the Tick–Bit substrate, established in companion VERSF papers and *imported* here.

Proposition 7.1 — Measurement as first passage

Within the Tick–Bit substrate, a measurement event is the first-passage commitment of the joint system + apparatus dynamics to a record-producing configuration. The realised outcome satisfies

$\text{outcome} = \arg\min_a T_a$,

where T_a *is the first-passage time to commitment along branch a, with rate law* $\lambda_a \propto |\psi_a|^2$ *consistent with §4.3.*

What is forced and what is imported

We foreground this calibration:

- **Forced by §3–6:** the pre-commitment kinematics, the amplitude representation, the form of the probability functional.
- **Imported from the Tick–Bit substrate:** the dynamical rate law $\lambda_a \propto |\psi_a|^2$, the first-passage mechanism, the apparatus microstate dynamics.
- **Required for consistency:** the imported rate law must agree with §4.3.

Proposition 7.2 — Substrate consistency

Let a candidate substrate dynamics produce a first-passage law $\lambda_a = f(|\psi_a|^2)$ *for some function f. The substrate is admissible only if* $f(x) \propto x$*.*

Proof. The kinematic Born rule (§4.3, §6.1) fixes $P(a) = |\psi_a|^2 / \Sigma_b |\psi_b|^2$. The substrate first-passage probability is $P_\text{subst}(a) = \lambda_a / \Sigma_b \lambda_b = f(|\psi_a|^2) / \Sigma_b f(|\psi_b|^2)$. For these to coincide on all admissible amplitude distributions — including those concentrated on a single component, distributed uniformly across N components, and arbitrary intermediate distributions — f must satisfy the constraint

$f(|\psi_a|^2) / \Sigma_b f(|\psi_b|^2) = |\psi_a|^2 / \Sigma_b |\psi_b|^2$

for all admissible $(|\psi_b|^2)_b \in \mathbb{R}_{\geq 0}^N$. Setting $\psi_b = 0$ for $b \neq a$, $b \neq c$ gives $f(x)/(f(x) + f(y)) = x/(x+y)$ for all $x, y \geq 0$, which forces $f(x)/f(y) = x/y$, i.e. $f(x) = \kappa x$ for some constant $\kappa > 0$. Normalisation fixes κ; without loss of generality $\kappa = 1$. ■

This rules out alternative rate laws (e.g. $f(x) = x^2$, $f(x) = \sqrt{x}$, $f(x) = \log(1+x)$) and fixes the $|\psi|^2$ rate as the unique admissible substrate law. The agreement of the Tick–Bit substrate with §4.3 is therefore *not* a coincidence but a structural consequence of the kinematic Born rule plus the requirement that the substrate reproduce it. Other substrates satisfying admissibility must produce the same rate law.

Consequences (within the substrate-augmented framework)

- **No collapse postulate.** The transition from amplitudes to outcomes is the resolution of a competitive first-passage process, not an extra dynamical law.

- **No fundamental randomness.** Probabilities arise from epistemic ignorance over the apparatus microstate; the rate law is fixed by §4.3 plus substrate consistency (Prop 7.2).
- **No measurement / unitary divide.** The same dynamics governs both regimes; what differs is whether an irreversible record is produced.

This is what we mean when we say the Born rule is *realised* on an admissible substrate rather than postulated. The realisation is substrate-dependent; the kinematic skeleton is not; and the substrate-consistency proposition ensures that any admissible substrate produces the same rate law.

---

8. Time and Irreversibility (forward pointer)

Two further results in the VERSF programme close the loop. The treatment here is *brief and forward-pointing*; the full arguments are developed in companion proto-time and emergent-Lorentz papers, to which this section refers.

- *Distinguishability ordering.* An ordering of physical states is well-defined only when those states are mutually distinguishable. Time is therefore not a background parameter but an *induced ordering* on commitment events.
- *Finitude and irreversibility.* Irreversibility requires finitude: an unbounded distinguishability budget would permit perfect undoing of any record.

Forward-pointer 8.1

Physical time is the partial order on the set of irreversible commitment events, equipped with a metric induced by distinguishability flow on the substrate.

We mark this as a *forward pointer* rather than a result of the present paper. The propositional content — partial order on commitment events — is a corollary of (A2) plus distinguishability ordering. The metric structure — which is the substantive claim — is established in the companion proto-time paper, and the recovery of Lorentz invariance from substrate-level distinguishability flow is established in the companion emergent-Lorentz paper. We do not derive these here; we cite them.

This places time *downstream* of the admissibility structure rather than alongside it. The arrow of time is the arrow of commitment.

---

9. Synthesis: The Admissibility Closure Theorem

We can now state the central result, calibrated to what the body of the paper establishes.

Theorem 9.1 — Admissibility Closure (kinematic part)

Any physical substrate satisfying (A1)–(A4) and (P) necessarily exhibits:

1. *a finite, structured set of pre-commitment alternatives;*
2. *a reversible amplitude representation* $\psi \in \mathbb{C}^N$ *carrying a global U(1) phase, where the selection of* $\mathbb{C}$ *over* $\mathbb{R}$ *and* $\mathbb{H}$ *follows from (A3) tomographic locality and (A4b) per-component continuous symmetry;*
3. *a bilinear, phase-invariant probability functional reducing to the Born rule*

$p_a = |\langle\varphi(a) \mid \psi\rangle|^2 / \Sigma_b |\langle\varphi(b) \mid \psi\rangle|^2$.

Theorem 9.2 — Admissibility Closure (substrate-dependent part)

Any admissible substrate that additionally supplies a first-passage commitment dynamics — and whose rate law is therefore constrained by Proposition 7.2 to satisfy $\lambda_a \propto |\psi_a|^2$ *— exhibits:*

4. *measurement as the realisation of that first-passage process, producing irreversible records;*
5. *physical time as the order induced on those records.*

Interpretation

The standard quantum postulates — Hilbert space, the Born rule, the projection postulate — are not independent assumptions at the kinematic level. They are facets of one admissibility constraint, and any universe sustaining stable irreversible records is forced into this kinematic architecture (Theorem 9.1). The dynamical realisation of measurement (Theorem 9.2) is substrate-dependent, but the substrate-consistency proposition (7.2) shows that any admissible substrate must reproduce the same rate law: the kinematic Born rule constrains the dynamical realisation rather than merely cohabiting with it.

The closure claim is therefore precise:

Quantum kinematics is forced by (A1)–(A4) and (P) modulo a stated field-selection step. Quantum measurement dynamics is forced by (A1)–(A4), (P), and the existence of a first-passage substrate, with the rate law uniquely determined by the kinematic skeleton.

This is weaker than "quantum mechanics is the unique admissible architecture" full stop — the field-selection step and the existence of a first-passage substrate are conditions, not consequences — and stronger than "quantum mechanics is one of many possible admissible architectures." It is what the argument actually delivers.

---

10. Conclusion

We summarise:

1. **Finite distinguishability** is necessary for stable facts.
2. **Irreversible commitment** plus the **channel-uniqueness premise** define what a measurement *is* and exclude commitment-free distinguishability change.
3. A **set-level reversible pre-commitment domain** is forced by (A1), (A2), and (P).
4. A **continuous representation-level symmetry** is forced by (A4a).
5. The **division algebra** carrying the representation is finite-dimensional, associative, continuous, and division-closed, hence (by Frobenius) $\mathbb{R}$, $\mathbb{C}$, or $\mathbb{H}$; the algebra is fixed to $\mathbb{C}$ by (A3) tomographic locality and (A4b) per-component continuous symmetry.
6. **Compositional consistency** forces the probability functional to be jointly bilinear and rank-one — the Born rule — in a single derivation.
7. **Measurement** is the first-passage realisation of this structure on the Tick–Bit substrate, with the rate law uniquely determined by Proposition 7.2.
8. **Time** is the order on commitment events (forward pointer to companion papers for metric structure).

Final statement (calibrated)

Quantum mechanics emerges as the unique admissible kinematic architecture under the stated constraints, modulo the field-selection step traced to tomographic locality and per-component continuous symmetry. Its measurement dynamics is the first-passage realisation of that kinematics on an admissible substrate, with the substrate's rate law uniquely fixed by kinematic consistency.

This is a strong claim — strong enough to make the standard postulates derivable rather than independent — and it is calibrated so the body of the paper actually earns it.

Appendix A — Bilinearity and rank-one factorisation, full calculation

We provide the explicit calculation underlying Theorem 4.3, in two parts: (A.1) the suppression of (n,n)-order terms with $n \geq 2$ by universal factorisation; (A.2) the rank-one closure of the surviving (1,1) kernel.

A.1 Suppression of higher-order terms

Let $P : \mathbb{C}^N \to \mathbb{R}_{\geq 0}$ be continuous, U(1)-phase-invariant, and complex-conjugation-invariant. Phase invariance restricts the formal series expansion of P around $\psi = 0$ to terms of the form

$P(\psi) = \Sigma_{n\geq 1} Q_n(\psi, \psi^*)$, where Q_n is a (conjugate-)symmetric form of bidegree (n, n) — i.e. degree n in ψ and degree n in ψ^* separately.

Now consider a product state $\psi = \psi^{(1)} \otimes \psi^{(2)} \in \mathbb{C}^{N1} \otimes \mathbb{C}^{N2}$. The components are $\psi_{ij} = \psi^{(1)}_i \cdot \psi^{(2)}_j$. Compositional consistency (A3) requires

$$P(\psi^{(1)} \otimes \psi^{(2)}) = P^{(1)}(\psi^{(1)}) \cdot P^{(2)}(\psi^{(2)})$$

for all admissible $\psi^{(1)}$, $\psi^{(2)}$ (after restricting to the appropriate marginal index).

Substitute the series expansions on both sides:

$$\Sigma_n Q_n(\psi^{(1)} \otimes \psi^{(2)}, \psi^{(1)*} \otimes \psi^{(2)*}) = (\Sigma_m Q_m^{(1)}) \cdot (\Sigma_k Q_k^{(2)}).$$

Compare both sides as formal series in the rescaled amplitudes $(\lambda\psi^{(1)}, \mu\psi^{(2)})$ with $\lambda, \mu \in \mathbb{R}_{>0}$. The left-hand side scales as $\Sigma_n |\lambda|^{2n} |\mu|^{2n} Q_n$; the right-hand side scales as $\Sigma_{m,k} |\lambda|^{2m} |\mu|^{2k} Q_m^{(1)} Q_k^{(2)}$.

Matching coefficients of $|\lambda|^{2a} |\mu|^{2b}$:

- For (a, b) with $a = b = n$: Q_n on the joint must equal $\Sigma_{m+k=n?\ no,\ m=k=n} Q_n^{(1)} Q_n^{(2)}$ — but more precisely, since the right-hand side requires equal scaling in λ and μ to match the left, we need $a = b$ for the scaling to balance, and within that diagonal: only the $m = k = a = b$ term contributes.

Hence at each n,

$$Q_n(\psi^{(1)} \otimes \psi^{(2)}, \cdot) = Q_n^{(1)}(\psi^{(1)}, \cdot) \cdot Q_n^{(2)}(\psi^{(2)}, \cdot).$$

This is a tensor-factorisation constraint *separately at each order n*. Now apply normalisation. Probability conservation $\Sigma_{\text{outcomes}} P_{\text{outcome}}(\psi) = 1$ across an outcome partition requires that *some* fixed-degree combination of the Q_n integrate to the constant 1 over the outcome set, for all ψ. Combined with the tensor factorisation at each order, this forces a Cauchy-type multiplicativity constraint on the diagonal pieces: defining $f(x) := \Sigma_{\text{outcomes}}\ Q_1$-component on a fixed ray $x = |\psi|^2$ — the contribution of Q_n on a ray scales as x^n, so $\Sigma_n c_n x^n$ must factorise across products $x = x^{(1)}x^{(2)}$, giving the functional equation

$$\Sigma_n c_n (xy)^n = (\Sigma_n c_n x^n)(\Sigma_n c_n y^n)$$

for all $x, y \geq 0$. The continuous solutions of this equation are $c_n = \delta_{n,1}$ (with $c_1 = 1$ by normalisation), or the trivial $c_n = 0$ case. Hence $Q_n = 0$ for all $n \geq 2$.

This is the suppression result: only the (1,1)-order term survives.

A.2 Rank-one closure of the surviving kernel

The surviving term is

$P(a \mid \psi) = \Sigma_{ij}\ \psi_i K_{ij}(a)\ \psi_j^* = \langle \psi \mid K(a) \mid \psi \rangle$,

with $K(a)$ Hermitian (by reality of P) and positive semidefinite (by $P \geq 0$).

Apply factorisation again at the (1,1) level on a product state:

$\langle \psi^{(1)} \otimes \psi^{(2)} \mid K(a, b) \mid \psi^{(1)} \otimes \psi^{(2)} \rangle = \langle \psi^{(1)} \mid K^{(1)}(a) \mid \psi^{(1)} \rangle \cdot \langle \psi^{(2)} \mid K^{(2)}(b) \mid \psi^{(2)} \rangle$

for all $\psi^{(1)}$, $\psi^{(2)}$. The left-hand side equals $\langle \psi^{(1)} \mid A \mid \psi^{(1)} \rangle$ for the partial-trace-like operator A defined by the action on $\psi^{(1)}$ at fixed $\psi^{(2)}$, and similarly with roles reversed. Universality over all $\psi^{(2)}$ forces the joint kernel $K(a, b)$ to have the tensor structure $K^{(1)}(a) \otimes K^{(2)}(b)$.

A Hermitian positive operator K such that *every* product-state expectation factorises as a product of marginal expectations, *for all* product states, is rank-one. This is a standard result in operator factorisation theory (see Bhatia 1997 §IV.4 on tensor-product factorisation of positive operators, and Nielsen–Chuang 2010 §2.5 for the equivalent statement in the language of pure-state separability of density operators). We give the elementary argument here for completeness. (Proof: if K had rank ≥ 2, there would exist orthogonal vectors u, v such that $\langle u \mid K \mid u \rangle > 0$ and $\langle v \mid K \mid v \rangle > 0$; on the product state $(u + v) \otimes (u + v)$ the expectation is $\langle u|K|u \rangle + \langle v|K|v \rangle + 2\,\mathrm{Re}\,\langle u|K|v \rangle$, which generically does not factorise as $\langle u{+}v \mid K^{(1)} \mid u{+}v \rangle \cdot \langle u{+}v \mid K^{(2)} \mid u{+}v \rangle$ unless $K^{(1)}$ and $K^{(2)}$ are themselves rank-one. Universality in N forces this at every system size.)

Hence $K(a) = |\varphi(a)\rangle\langle\varphi(a)|$ for some $\varphi(a) \in \mathbb{C}^N$, and

$P(a \mid \psi) = |\langle \varphi(a) \mid \psi \rangle|^2$.

Normalisation $\Sigma_a P(a \mid \psi) = 1$ for all ψ requires the $\varphi(a)$ to form a POVM resolution $\Sigma_a |\varphi(a)\rangle\langle\varphi(a)| = \mathbb{1}_N$, recovering the standard form. ■

References

Araki, H. (1980). On a characterization of the state space of quantum mechanics. *Communications in Mathematical Physics* 75, 1–24.

Baez, J. C. (2002). The octonions. *Bulletin of the American Mathematical Society* 39, 145–205.

Baez, J. C. (2012). Division algebras and quantum theory. *Foundations of Physics* 42, 819–855.

Barrett, J. (2007). Information processing in generalized probabilistic theories. *Physical Review A* 75, 032304.

Bhatia, R. (1997). *Matrix Analysis*. Graduate Texts in Mathematics 169. Springer.

Busch, P. (2003). Quantum states and generalized observables: a simple proof of Gleason's theorem. *Physical Review Letters* 91, 120403.

Chiribella, G., D'Ariano, G. M., & Perinotti, P. (2011). Informational derivation of quantum theory. *Physical Review A* 84, 012311.

Frobenius, F. G. (1878). Über lineare Substitutionen und bilineare Formen. *Journal für die reine und angewandte Mathematik* 84, 1–63.

Gleason, A. M. (1957). Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics* 6, 885–893.

Hardy, L. (2001). Quantum theory from five reasonable axioms. arXiv:quant-ph/0101012.

Masanes, L., & Müller, M. P. (2011). A derivation of quantum theory from physical requirements. *New Journal of Physics* 13, 063001.

Mohrhoff, U. (2004). Probabilities from envariance? *International Journal of Quantum Information* 2, 221–229.

Nielsen, M. A., & Chuang, I. L. (2010). *Quantum Computation and Quantum Information* (10th anniversary edition). Cambridge University Press.

Schlosshauer, M., & Fine, A. (2005). On Zurek's derivation of the Born rule. *Foundations of Physics* 35, 197–213.

Stueckelberg, E. C. G. (1960). Quantum theory in real Hilbert space. *Helvetica Physica Acta* 33, 727–752.

Wootters, W. K. (1990). Local accessibility of quantum states. In *Complexity, Entropy, and the Physics of Information*, ed. W. H. Zurek, 39–46. Addison-Wesley.

Zurek, W. H. (2005). Probabilities from entanglement, Born's rule p_k = |ψ_k|² from envariance. *Physical Review A* 71, 052105.